

Wikiprint Book

Title: System usage

Subject: DEEP - Public/User_Guide/DEEP-EST_CM

Version: 15

Date: 15.05.2024 08:24:47

Table of Contents

System usage	3
Filesystems and local storage	3
Multi-node Jobs	3

System usage

The DEEP-EST Cluster Module (CM) can be used through the SLURM based batch system that is also used for (most of) the Software Development Vehicles (SDV). You can request CM cluster nodes (dp-cn[01-50]) on the SDV with an interactive session like this:

```
srun --partition=dp-cn -N 4 -n 2 --pty /bin/bash -i
srun ./hello_cluster
Hello World from processor dp-cn15, rank 2 out of 8
Hello World from processor dp-cn15, rank 3 out of 8
Hello World from processor dp-cn17, rank 6 out of 8
Hello World from processor dp-cn17, rank 7 out of 8
Hello World from processor dp-cn14, rank 0 out of 8
Hello World from processor dp-cn16, rank 4 out of 8
Hello World from processor dp-cn14, rank 1 out of 8
Hello World from processor dp-cn16, rank 5 out of 8
```

When using a batch script, you have to adapt the partition option within your script: `--partition=dp-cn`

Filesystems and local storage

The home filesystem on the DEEP-EST Cluster Module is provided via GPFS/NFS and hence the same as on (most of) the remaining compute nodes. The local storage system of the CM running BeeGFS is available at

```
/work
```

There is a gateway being used to bridge between the Infiniband EDR used for the CM and the 40 GbE network the file servers are connected to.

This is NOT the same storage being used on the DEEP-ER SDV system. Both, the DEEP-EST prototype system and the DEEP-ER SDV have their own local storage.

It's possible to access the local storage of the DEEP-ER SDV (`/sdv-work`), but you have to keep in mind that the file servers of that storage can just be accessed through 1 GbE ! Hence, it should not be used for performance relevant applications since it is much slower than the DEEP-EST local storages mounted to `/work`.

There is also some node local storage available for the DEEP-EST Cluster nodes mounted to `/scratch` on each node (about 380 GB with XFS). Remember that this scratch is not persistent and **will be cleaned after your job has finished !**

Multi-node Jobs

The latest `pscom` version used in ParaStation MPI provides support for the Infiniband interconnect used in the DEEP-EST Cluster Module. Hence, loading the most recent ParaStation module will be enough to run multi-node MPI jobs over Infiniband:

```
module load ParaStationMPI
```

For using Cluster nodes in heterogenous jobs together with DAM nodes, please see

https://deeptrac.zam.kfa-juelich.de:8443/trac/wiki/Public/User_Guide/Batch_system#Heterogeneousjobs.